Simulating High-Speed Flows with NASA FUN3D

Gabriel Nastac
gabriel.c.nastac@nasa.gov
Computational AeroSciences Branch
NASA Langley Research Center

Overview

- Summary of FUN3D Capabilities
- High Performance Computing
  - Graphics Processing Units (GPUs)
- Retropropulsion Trajectory Simulations
- Automated Unstructured Grid Adaptation

Fully Unstructured Navier-Stokes 3D (FUN3D)

https://fun3d.larc.nasa.gov

- **Solves the time-dependent compressible Navier-Stokes equations in fully implicit form on unstructured grids with thermochemical nonequilibrium and assorted turbulence treatments**
- Combinations of rigid / overset / deforming grids and 6DOF
- Aeroelastic modeling using mode shapes, full FEM, etc.
- Constrained / multipoint adjoint-based design
- Grid adaptation
- Capabilities fully integrated, online documentation, training videos, tutorials
- **Lightweight abstraction over C++ enables performance across NVIDIA / AMD / Intel GPUs and multicore CPUs**

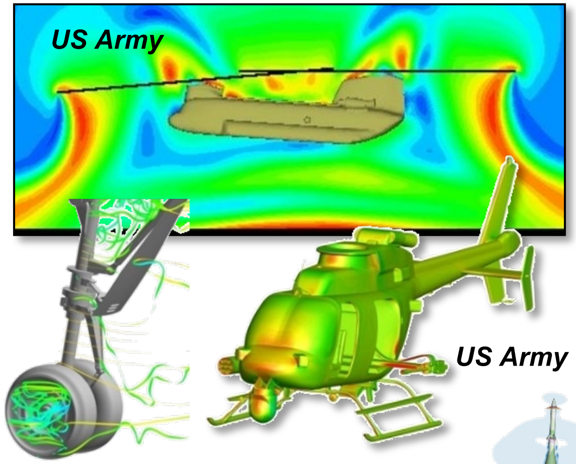

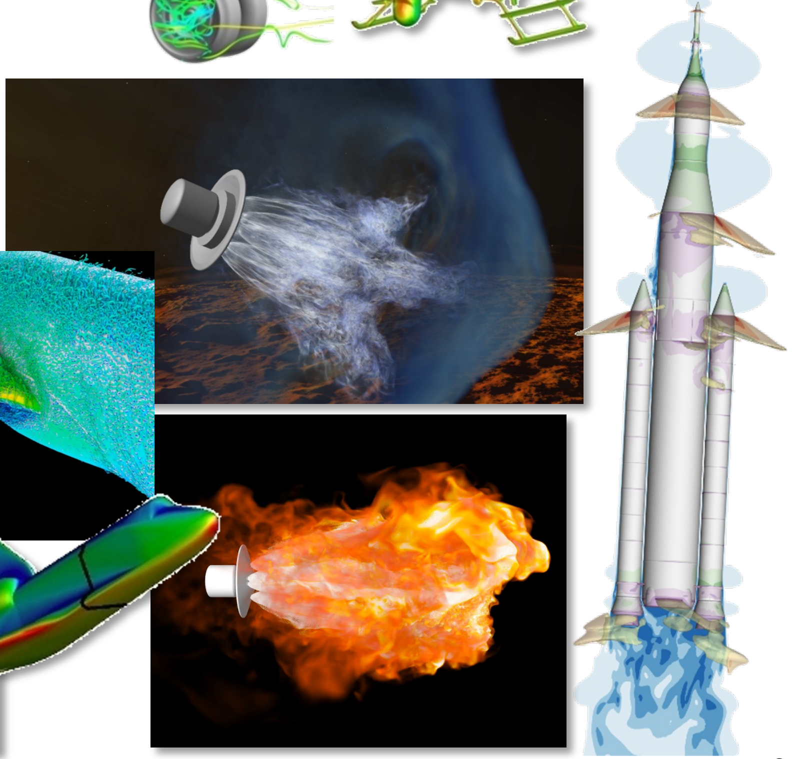

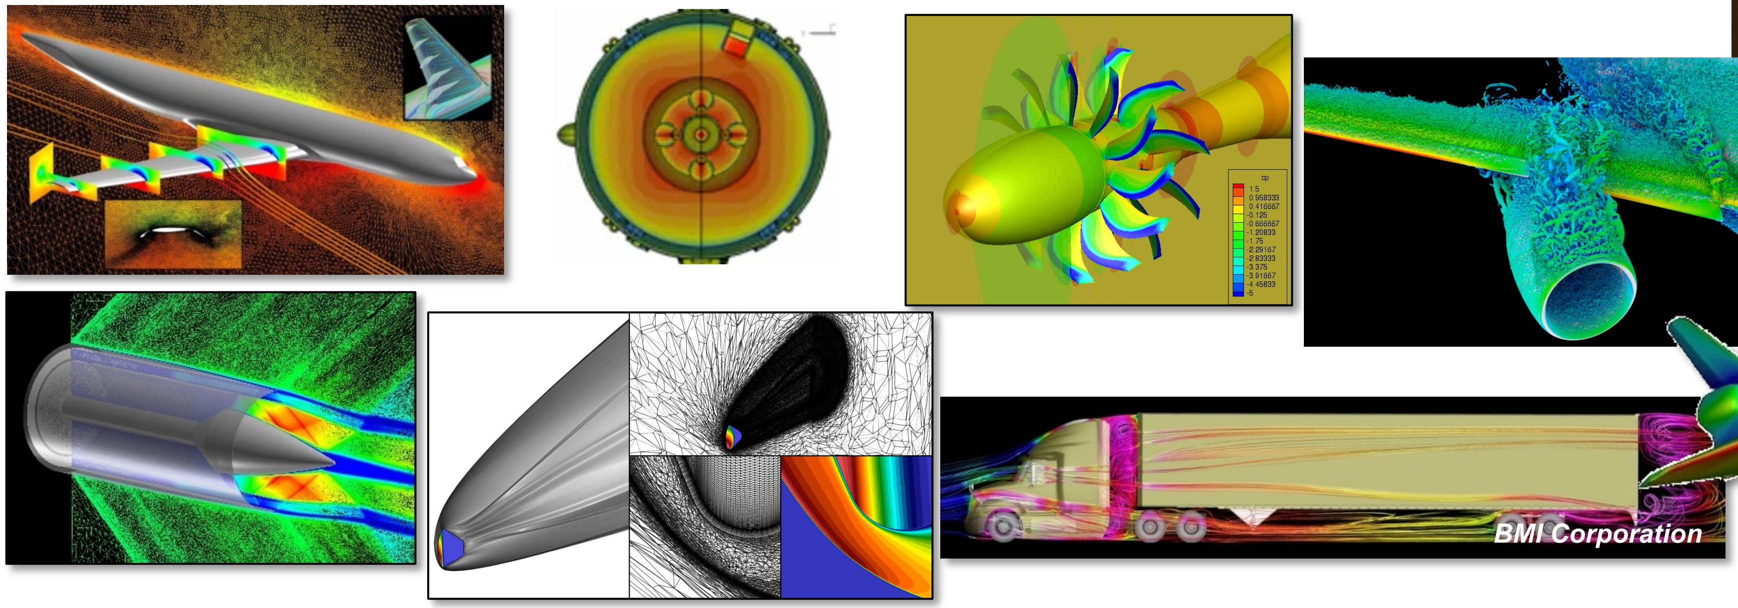

The CFD Vision 2030 Study

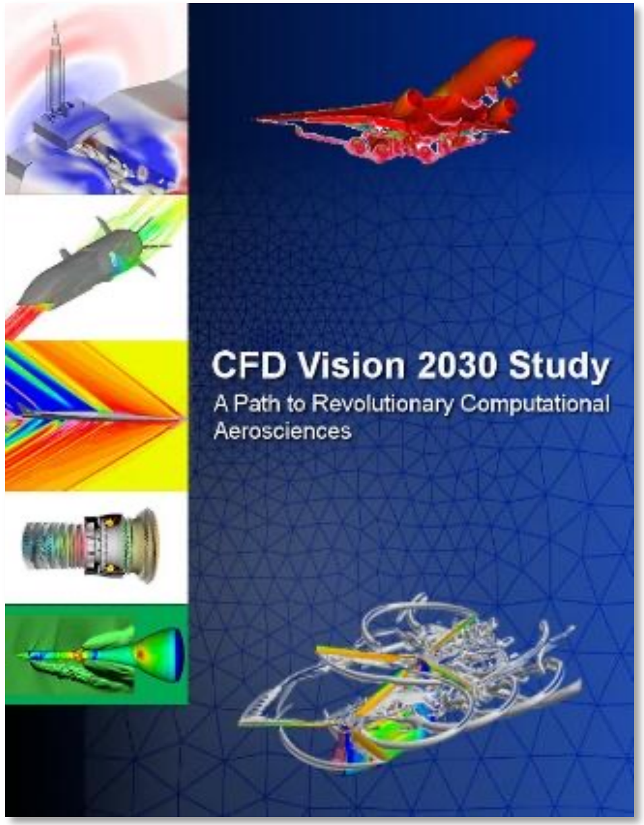

NASA/CR-2014-218178

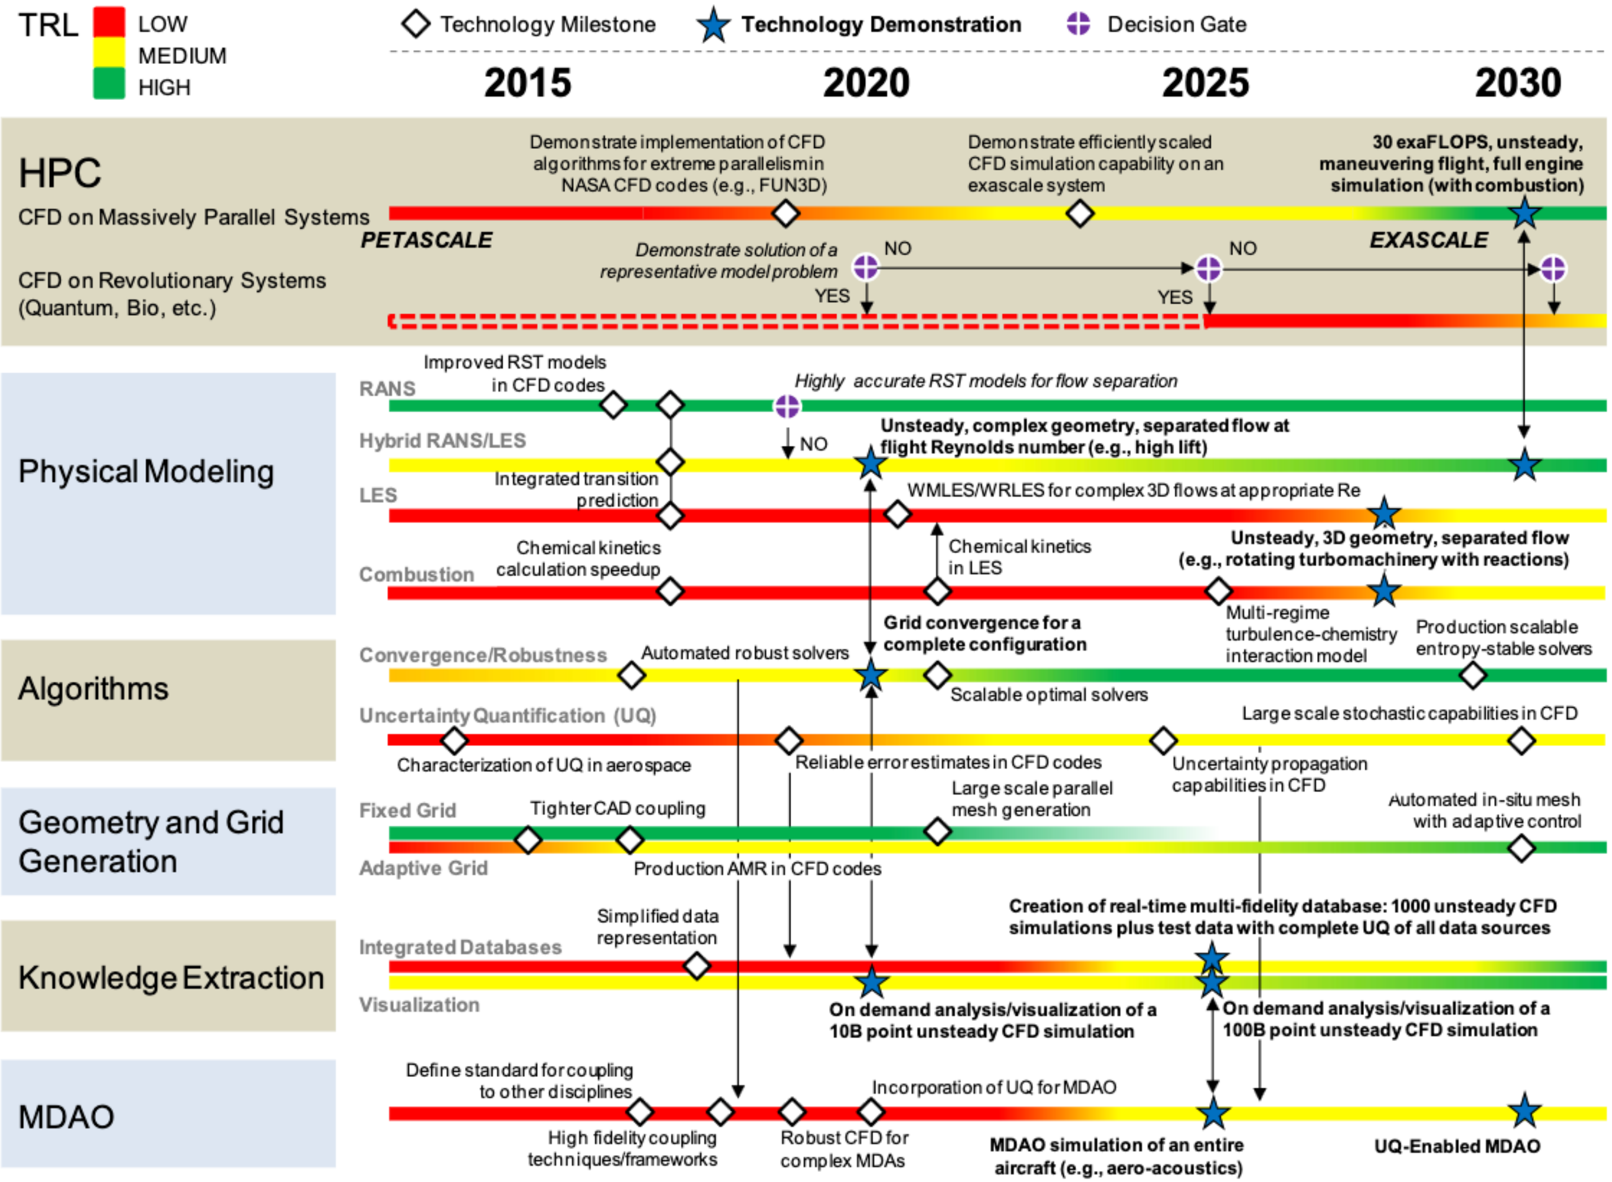

A Wake-Up Call for US Supercomputing

https://commons.wikimedia.org/wiki/File:EarthSimulator.jpg; Manatee_tw; Earth Simulator in Japan JAMSTEC 2007; https://creativecommons.org/licenses/by-sa/2.0/

New York Times Front Page, April 20, 2002:

"Japanese Computer is World's Fastest, As U.S. Falls Back"

"A Japanese laboratory has built the world's fastest computer, a machine so powerful that it matches the raw processing power of the 20 fastest American computers combined…"

"For some American computer scientists, the arrival of the Japanese supercomputer evokes the type of alarm raised by the Soviet Union's Sputnik satellite in 1957."

Computational speed: 36 Teraflops, or 36 x 10^{12} operations per second

Summer 2022: The Exascale Era Has Arrived

With the latest GPUs, we can now hold the 2002 Japanese system in the palm of our hand

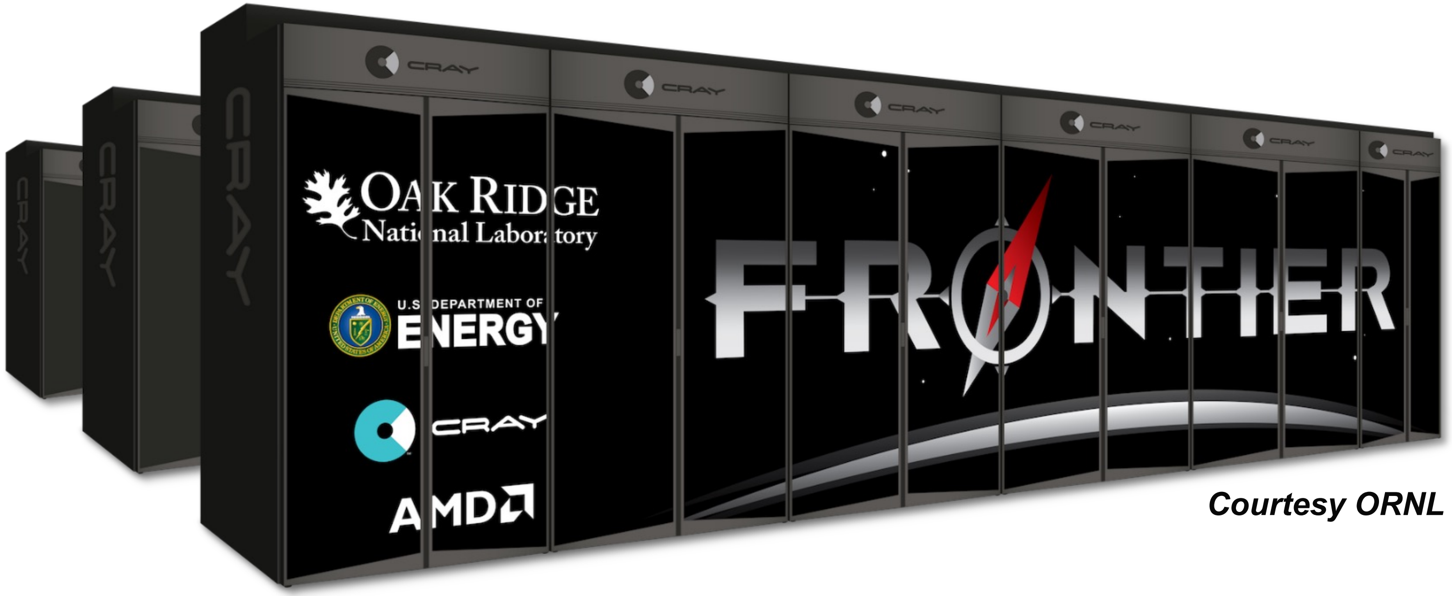

Courtesy ORNL

The Frontier system at Oak Ridge National Laboratory was certified as the world's first exascale system in Summer 2022

FUN3D simulations have been run across the entire system

Frontier first US Exascale computer

Multiple GPU per CPU drove energy efficiency

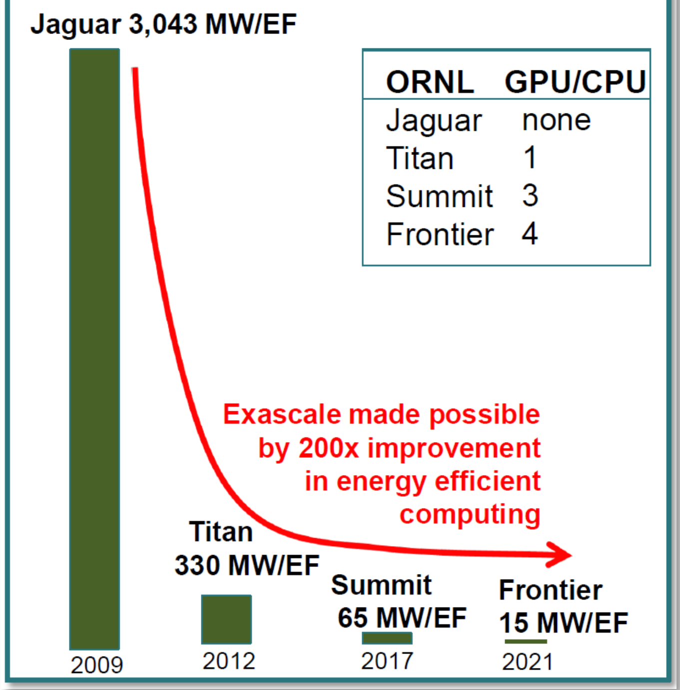

ORNL	GPU/CPU
Jaguar	none
Titan	1
Summit	3
Frontier	4

Courtesy Justin Whitt, ORNL

3000 MW ~ 200,000 Electric Cars on HW

15 MW ~ 1,000 Electric Cars on HW

High Performance Computing and GPUs

Global HPC Landscape

Rank	Organization	Name	Rmax (PF)	Installation Year
1	ORNL	Frontier	1200	2022
2	ANL	Aurora	585	2023
3	Microsoft	Eagle	561	2023
4	RIKEN	Fugaku	442	2020
5	EuroHPC	LUMI	380	2020
6	EuroHPC	Leonardo	239	2020
7	ORNL	Summit	150	2018
8	EuroHPC	MareNostrum	138	2023
9	NVIDIA	SuperPOD	121	2023
10	LLNL	Sierra	95	2018
85	NASA	Aitken	9	2021
98	DoD	Narwhal	8	2021

Upcoming US Systems in 2024
ANL Aurora (2000 PF)
LLNL El Capitan (2000 PF)

Architecture: CPU / GPU
PF: PetaFLOPS (10^{15} FLOPs)

Courtesy Top500.org

- Graphics Processing Units, or GPUs, enable next-generation performance with vastly reduced power requirements
- GPUs exploit orders of magnitude more parallelism than CPUs
- Requires application developers to retool existing algorithms and implementations

FUN3D GPU Strategy

- **Performance is paramount for GPU adoption: rearchitecting must prove cost-effective**
- Developed native lightweight abstraction[1] over NVIDIA CUDA C++
  - Enables efficient execution on NVIDIA / AMD / Intel GPUs and CPUs
- Hierarchical parallelism is essential for performance and is heavily leveraged
- **Like many large-scale science applications, FUN3D is primarily memory bound**

Architecture, CPU / GPU	Memory Bandwidth (GB/s)	Power (Watts)	Ratio
Intel Skylake 6148 Dual-Socket (40-core)	256	300	1
AMD EPYC 7762 Dual-Socket (128-core)	409.6	450	1.1
NVIDIA H100-NVL	7800	700	13.1
AMD MI300X	5218	750	8.2
Intel Data Center GPU Max 1550	3276.8	600	6.4

[1] Nastac, G., Walden, A., Wang, L., Nielsen, E., Liu, Y., Opgenorth, M., Orender, J. and Zubair, M., A Multi-Architecture Approach for Implicit Computational Fluid Dynamics on Unstructured Grids. AIAA Paper 2023-1226, 2023.

Retropropulsion Trajectory Simulations

Retropropulsion for Human Mars Exploration

- Human-scale Mars landers require new approaches to all phases of Entry, Descent, and Landing
- Cannot use heritage, low-L/D rigid capsules → deployable hypersonic decelerators or mid-L/D rigid aeroshells
- Cannot use parachutes → retropropulsion, from supersonic conditions to touchdown
- Limited understanding / numerous questions; physical testing is infeasible

	Viking	Pathfinder	MERs	Phoenix	MSL	InSight	M2020	Human-Scale Lander (Projected)
Entry Capsule (to scale)								
Diameter (m)	3.505	2.65	2.65	2.65	4.52	2.65	4.5	16 - 19
Entry Mass (t)	0.930	0.584	0.832	0.573	3.153	0.608	3.440	40 - 65
Parachute Diameter (m)	16.0	12.5	14.0	11.8	19.7	11.8	21.5	N/A
Parachute Deploy (Mach)	1.1	1.57	1.77	1.65	1.75	1.66	1.75	N/A
Landed Mass (t)	0.603	0.360	0.539	0.364	0.899	0.375	1.050	26 - 36
Landing Altitude (km)	-3.5	-2.5	-1.4	-4.1	-4.4	-2.6	-2.5	+/- 2.0
Landing Technology	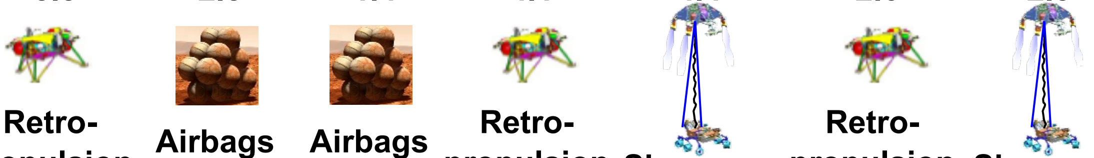 Retro-propulsion	Airbags	Airbags	Retro-propulsion	Skycrane	Retro-propulsion	Skycrane	Retro-propulsion

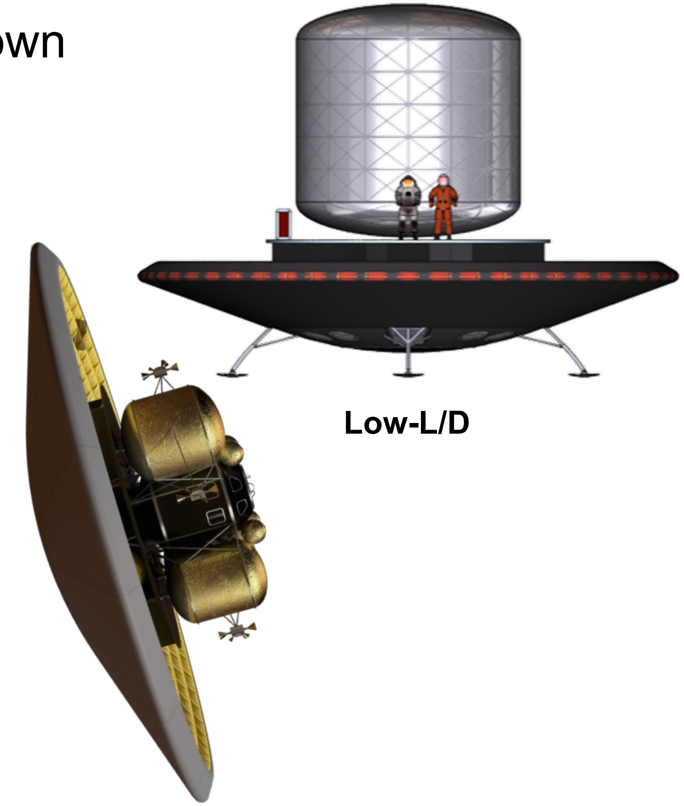

Steady progression of "in family" EDL → New EDL Paradigm

Summit Campaigns at Oak Ridge (2018-2021)

Static Simulations of Retropropulsion for Human-Scale Mars Landers

- LOX / CH_4 rocket engines in Martian CO_2 atmosphere
- 10 species / 19 reactions (15 PDEs); ~7 billion-element grids
- Game-changing computational performance: Two-day runs simulating seconds of real time on 16,000 NVIDIA V100 GPUs
  - Equivalent of several million CPU cores
- 90 GB of asynchronous I/O every 30 seconds for two days yields ~1 PB / run

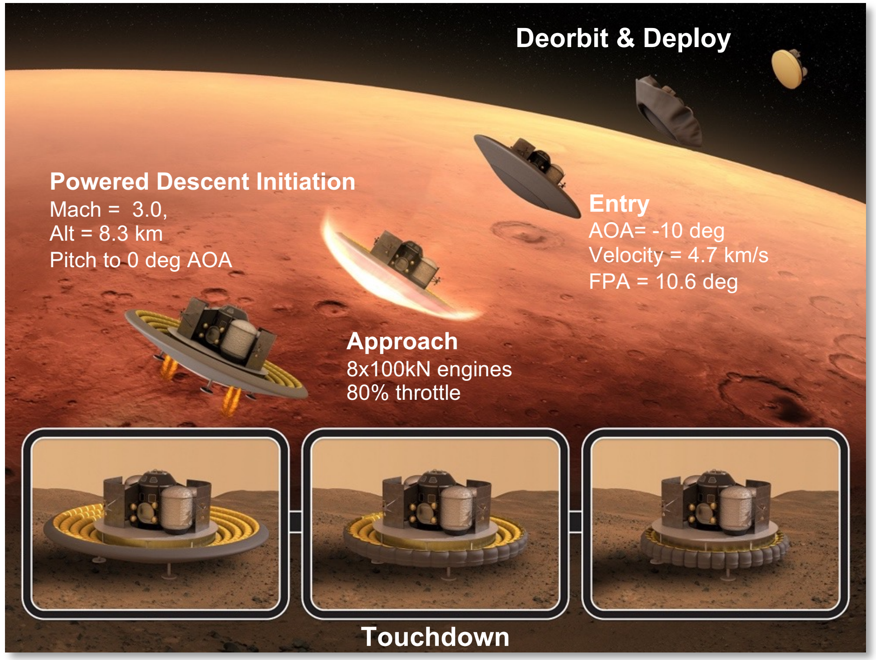

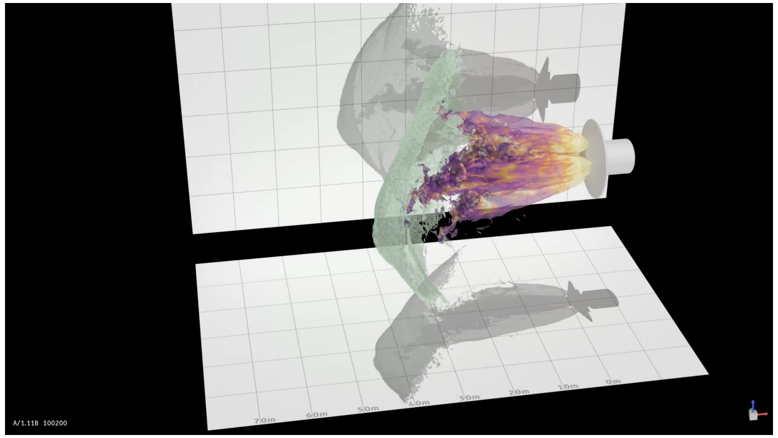

Mach 2.4: Dynamic Turbulent Flowfield

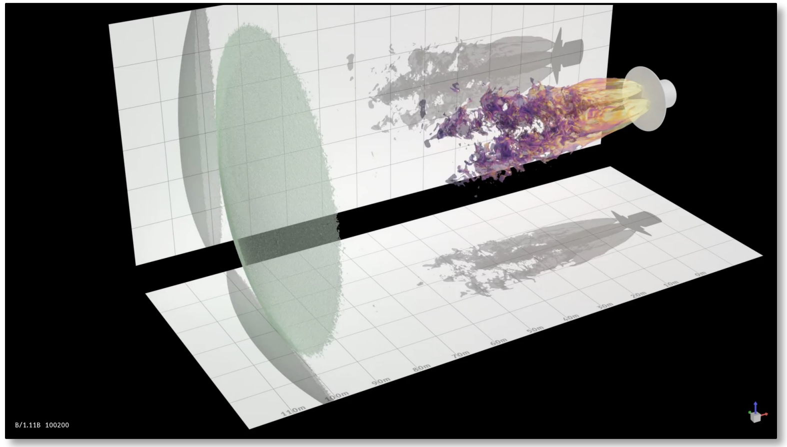

Mach 1.4: ~100 m Shock Standoff Distance

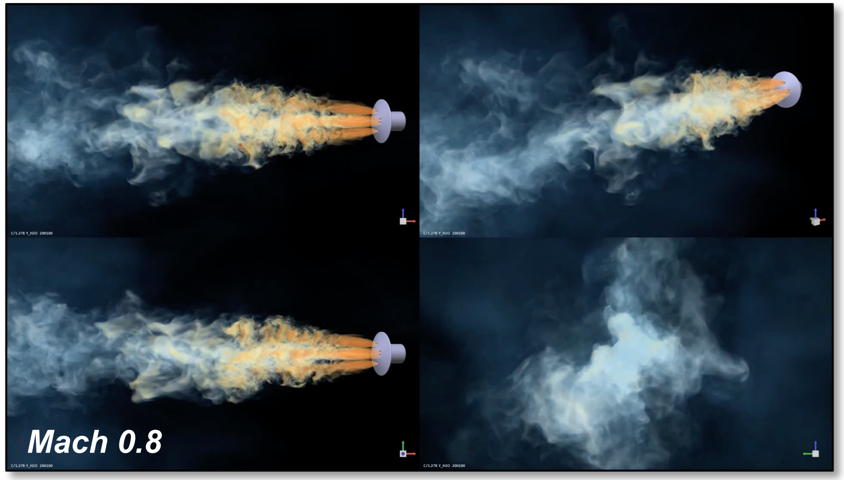

CFD-Based Trajectory Analysis

- Currently rely on Monte Carlo trajectory sims coupled with RANS-based databases
- However, many vehicles call for scale-resolving sims
  - Database cost increases dramatically
  - Dubious impact on uncertainties; hysteresis ignored
  - Some phenomena observed in static sims irrelevant in flight
- **CFD-based trajectory analysis is a grand-challenge problem recently posed to the aerospace CFD community**
  - **Augment static sims by "flying" high-fidelity, physics-based trajectories coupled with flight mechanics**
  - **Enables testing of flight mechanics algorithms in a high-fidelity simulation environment**

The new generation of exascale systems enables CFD-based trajectory analysis using scale-resolving methods

Typical CPU-Based Costs for Space Launch System Database Generation

Database	Code	Solutions (Grid Size)	Wall Time
Ascent	FUN3D	1,380 (60M)	2-4 weeks
Ascent	OVERFLOW	1,000 (500M)	2-3 months
F & M Wind Tunnel	FUN3D	600 (40M)	1 week
Booster Separation	FUN3D	13,780	3 months
Booster Separation	Cart3D	25,000	3 months

Derek Dalle, NASA ARC

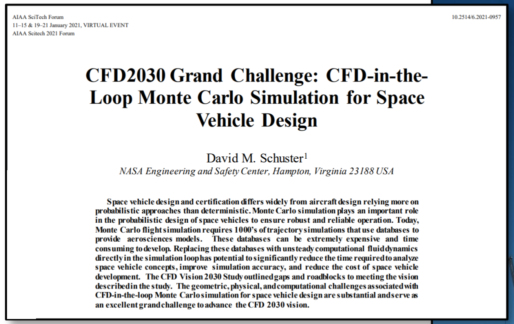

AIAA SciTech Forum
11–15 & 19–21 January 2021, VIRTUAL EVENT
AIAA Scitech 2021 Forum

10.2514/6.2021-0957

CFD2030 Grand Challenge: CFD-in-the-Loop Monte Carlo Simulation for Space Vehicle Design

David M. Schuster[1]
NASA Engineering and Safety Center, Hampton, Virginia 23188 USA

Space vehicle design and certification differs widely from aircraft design relying more on probabilistic approaches than deterministic. Monte Carlo simulation plays an important role in the probabilistic design of space vehicles to ensure robust and reliable operation. Today, Monte Carlo flight simulation requires 1000's of trajectory simulations that use databases to provide aerosciences models. These databases can be extremely expensive and time consuming to develop. Replacing these databases with unsteady computational fluid dynamics directly in the simulation loop has potential to significantly reduce the time required to analyze space vehicle concepts, improve simulation accuracy, and reduce the cost of space vehicle development. The CFD Vision 2030 Study outlined gaps and roadblocks to meeting the vision described in the study. The geometric, physical, and computational challenges associated with CFD-in-the-loop Monte Carlo simulation for space vehicle design are substantial and serve as an excellent grand challenge to advance the CFD 2030 vision.

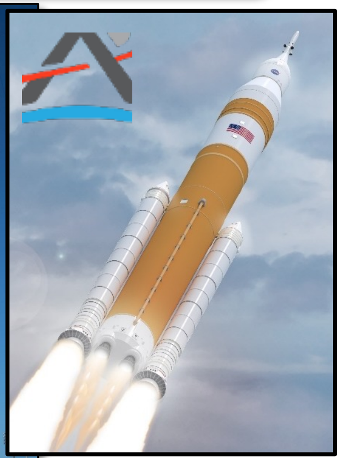

Source: NASA

Program to Optimize Simulated Trajectories II
(POST2)

- **Flight-validated, generalized, event-based, point-mass vehicle & trajectory simulation codebase**
  - 3 / 6 / Multi-DOF
  - Interfaces with user-provided multidisciplinary engineering models and flight software
  - Built-in trajectory optimization
- **Key Applications**
  - Statistical analysis of integrated performance
  - Orbital & atmospheric trajectory optimization and design
  - GN&C algorithm development & assessment
  - Off-nominal, faults, aborts, and margin analysis

Coupling FUN3D and POST2 provides far more capability than a traditional 6DOF implementation within the CFD solver

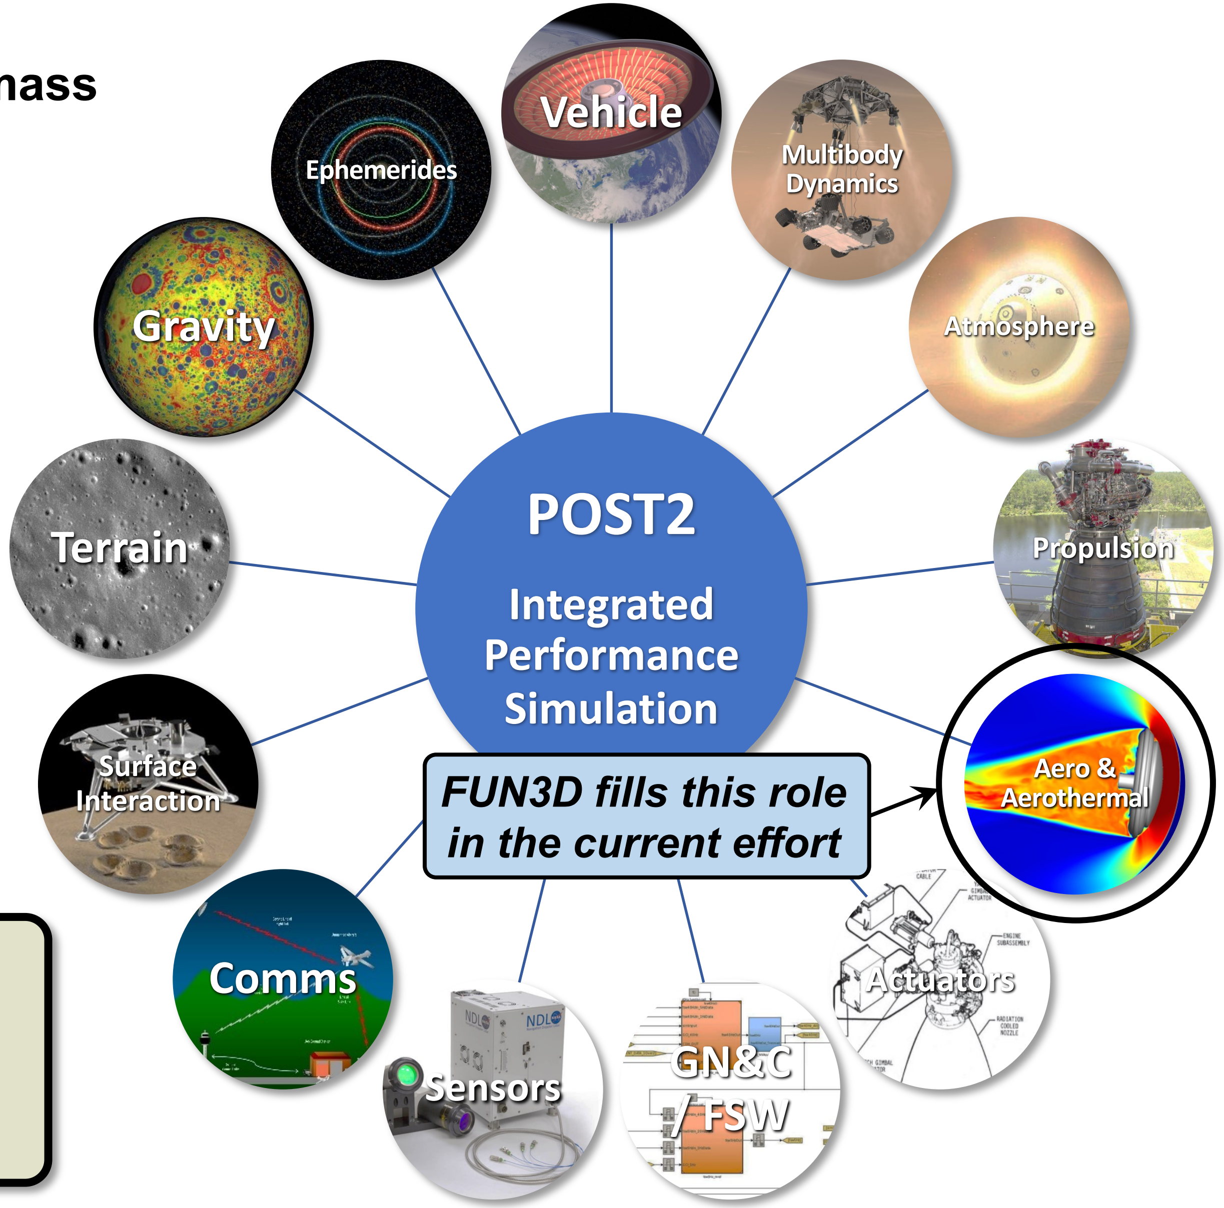

Example Trajectory Simulation

2022-2023 Oak Ridge Campaigns on Summit and Frontier

- Volume rendering of total temperature for 35-second trajectory (4 days of runtime)
- 7 billion elements
- 1-km spacing indicated on Martian surface
- Plumes grow substantially as vehicle decelerates from Mach 2.4 to Mach 0.8

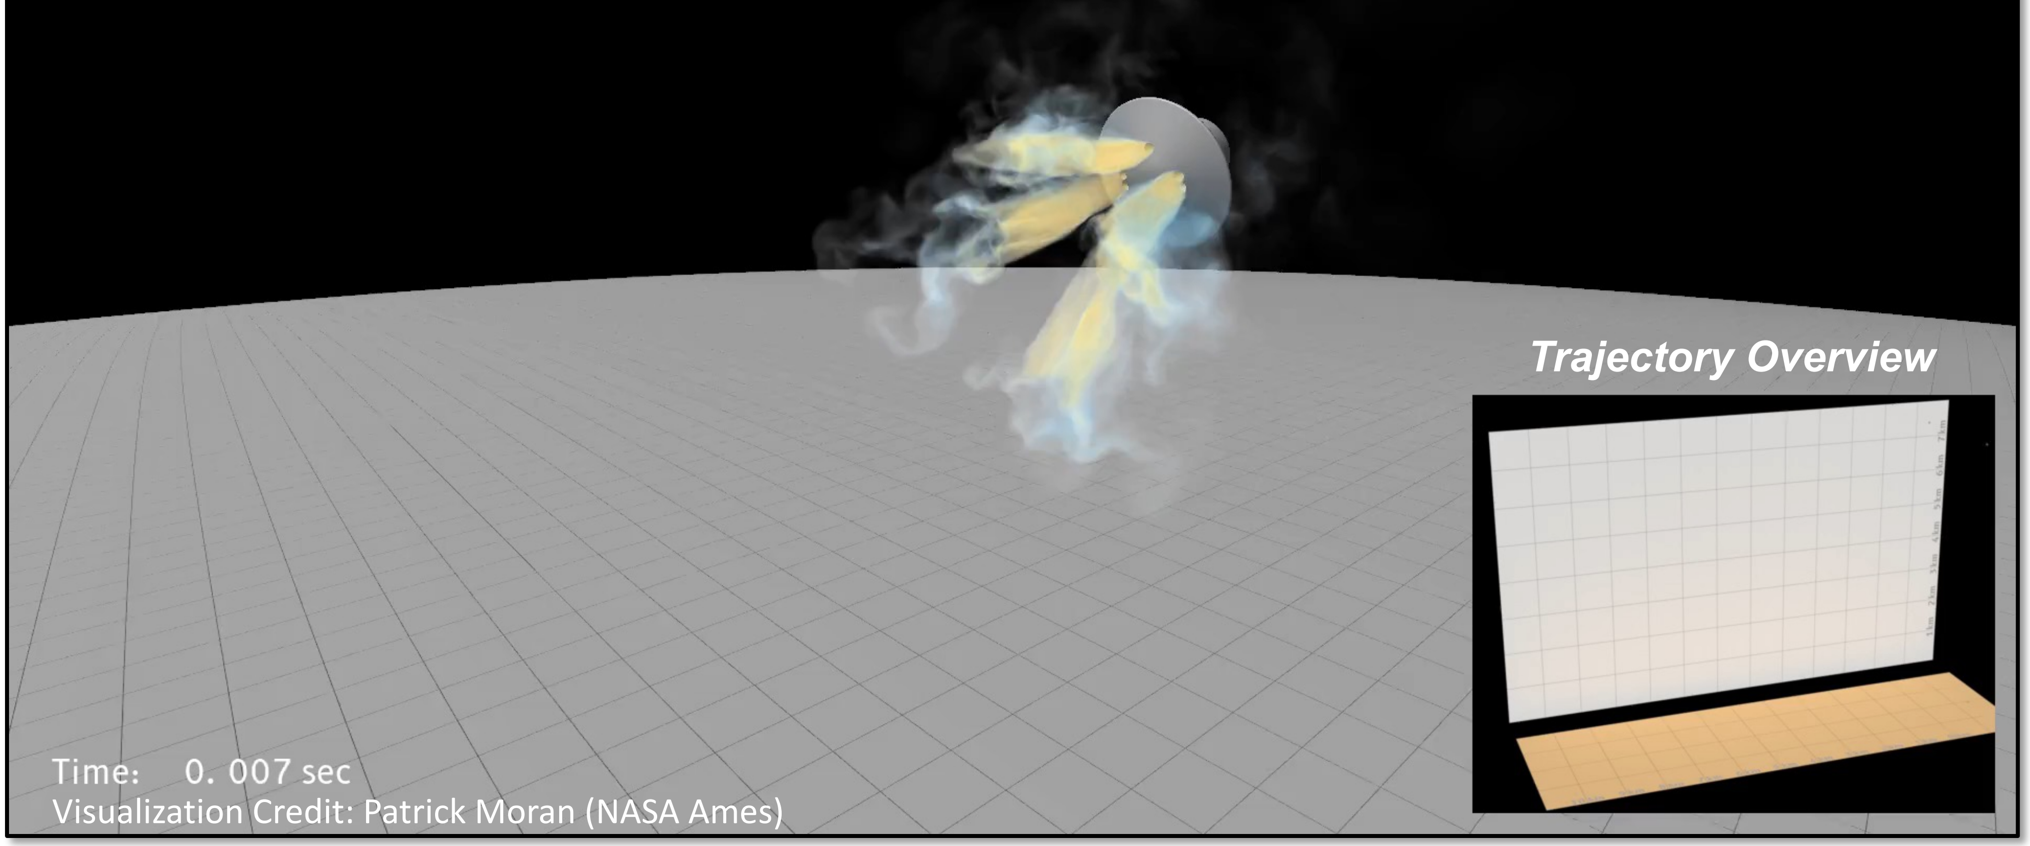

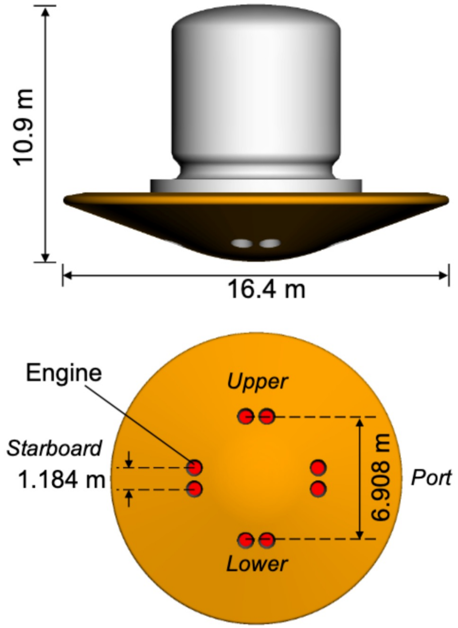

Automated Unstructured Grid Adaptation

Mixed-Element Unstructured Grid Adaptation

- Turbulent hypersonic flow over a basic 3D rocket geometry, 8M grid points at final cycle
- CFD cycle takes O(1 minute) on 1 8xA100 GPU node

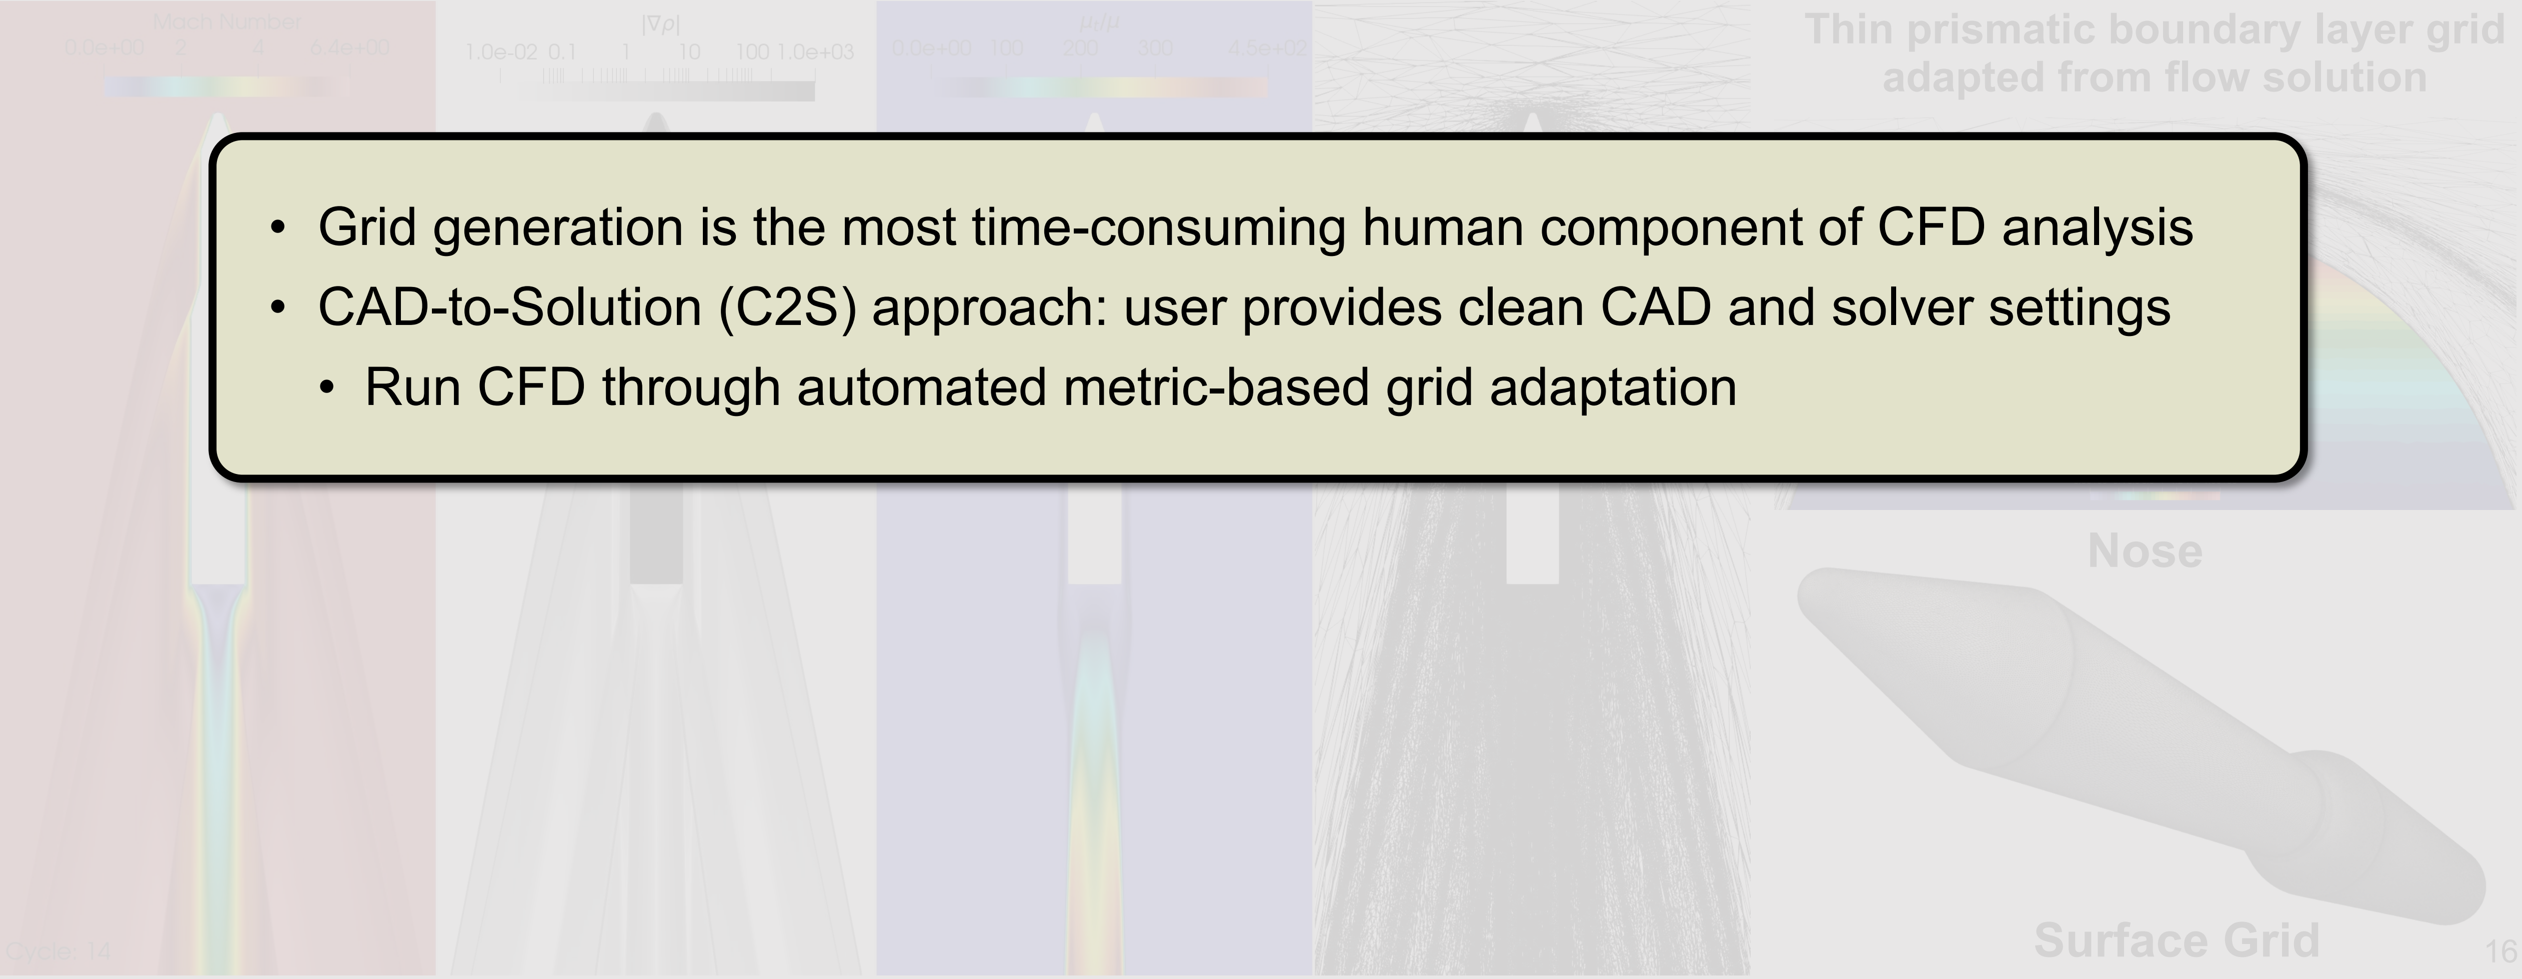

- Grid generation is the most time-consuming human component of CFD analysis
- CAD-to-Solution (C2S) approach: user provides clean CAD and solver settings
  - Run CFD through automated metric-based grid adaptation

Mixed-Element Unstructured Grid Adaptation

- **Turbulent hypersonic flow over a basic 3D rocket geometry, 8M grid points at final cycle**
- **CFD cycle takes O(1 minute) on 1 8xA100 GPU node**

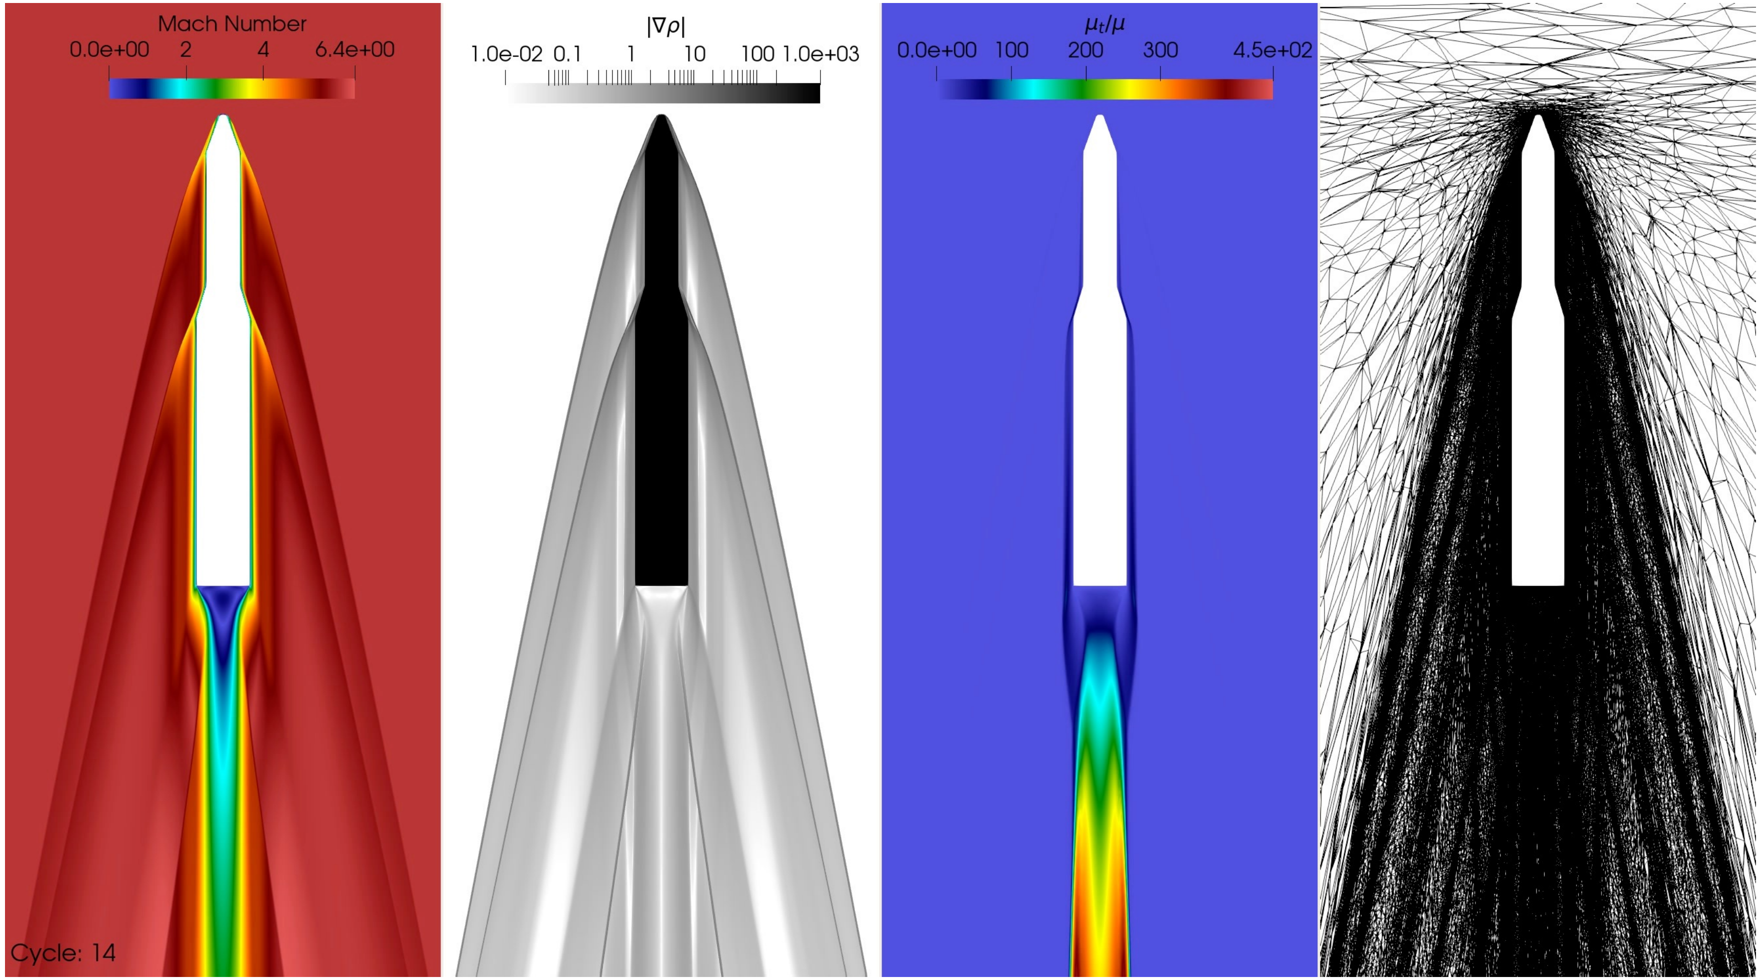

Thin prismatic boundary layer grid adapted from flow solution

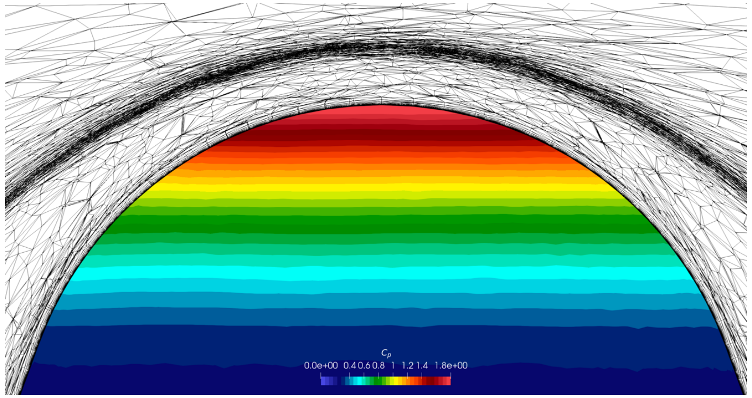

Nose

Surface Grid

Crew Exploration Vehicle

- $u_\infty = 4.6\frac{km}{s}, \alpha = 28°$
- Laminar flow, 5-species air
- CAD input; final grid is 4 million points and 19 million elements
- CFD cycle takes O(2 minutes) on 1 8xA100 GPU node

C_p

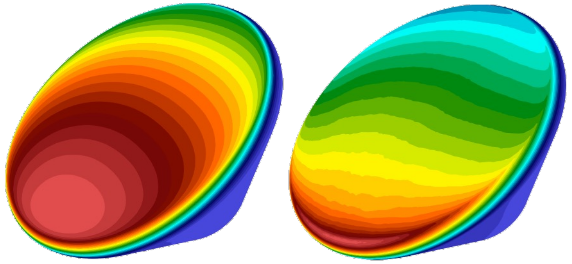

Heat Transfer

Surface Mesh

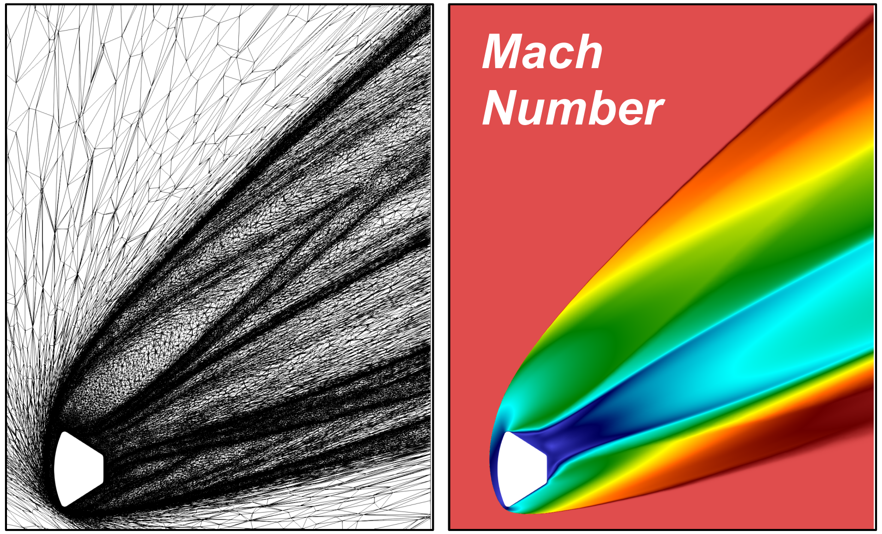

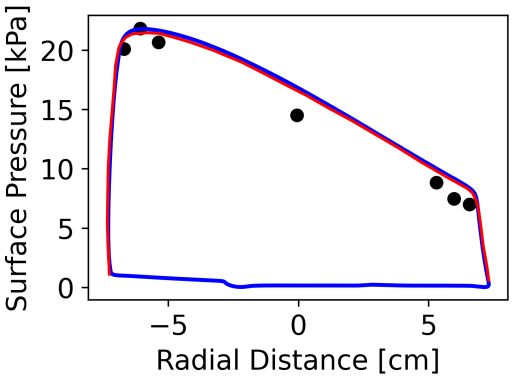

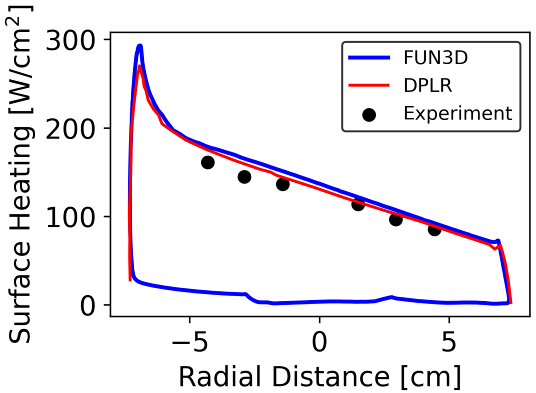

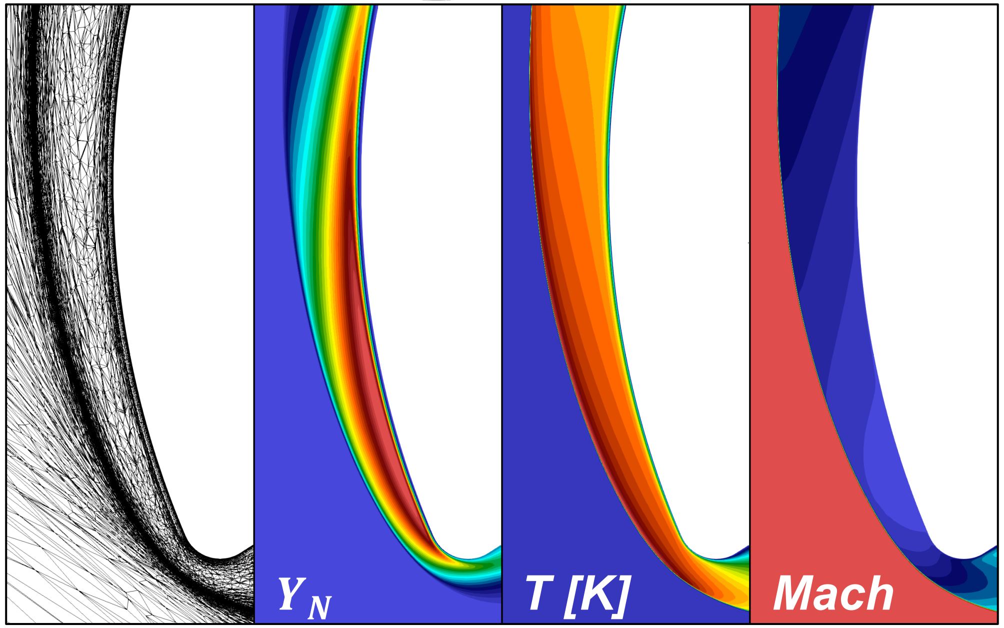

[1]Nastac, G., Tramel, R., & Nielsen, E., Improved Heat Transfer Prediction for High-Speed Flows over Blunt Bodies using Adaptive Mixed-Element Unstructured Grids, AIAA Paper 2022-0111, 2022

Sierra Space Dream Chaser

- Fixed prismatic boundary layer and surface grid, with tetrahedral adaptation using NASA *refine*
- Laminar flow, 5-species air
- $u_\infty = 5.0 \frac{km}{s}, \alpha = 40°$
- 15 cycles: 12 million points, 50 million elements for final grid
  - CFD cycle takes O(5 minutes) on 1 8xA100 GPU node

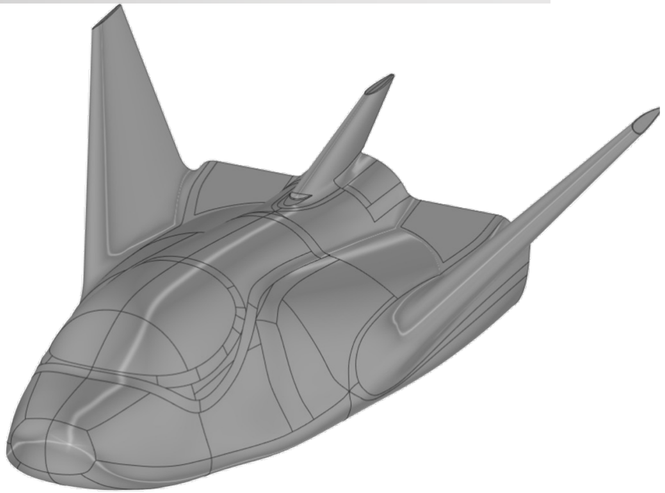

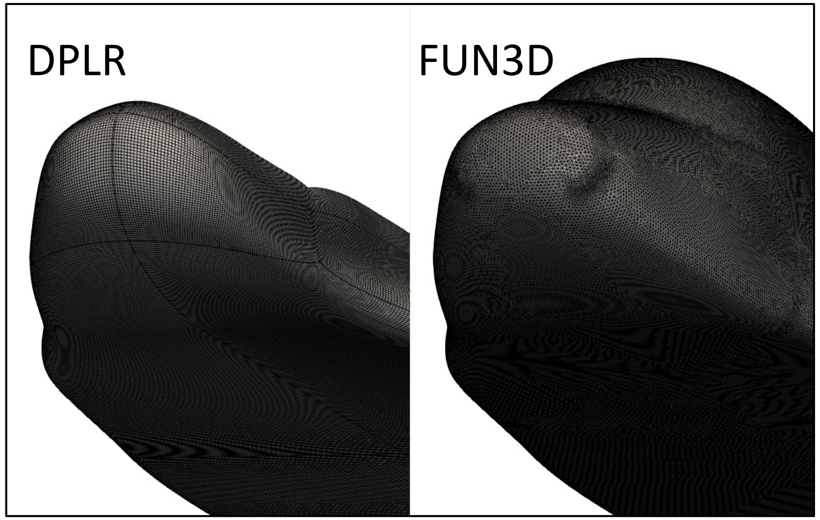

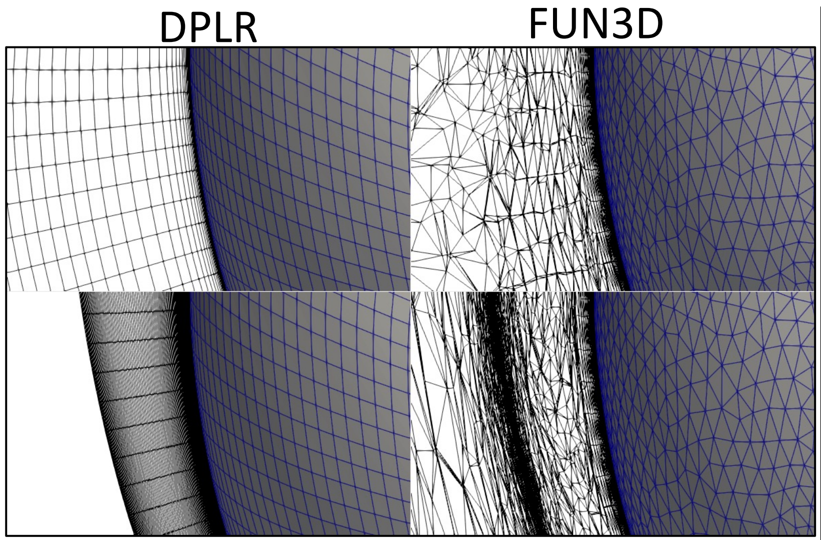

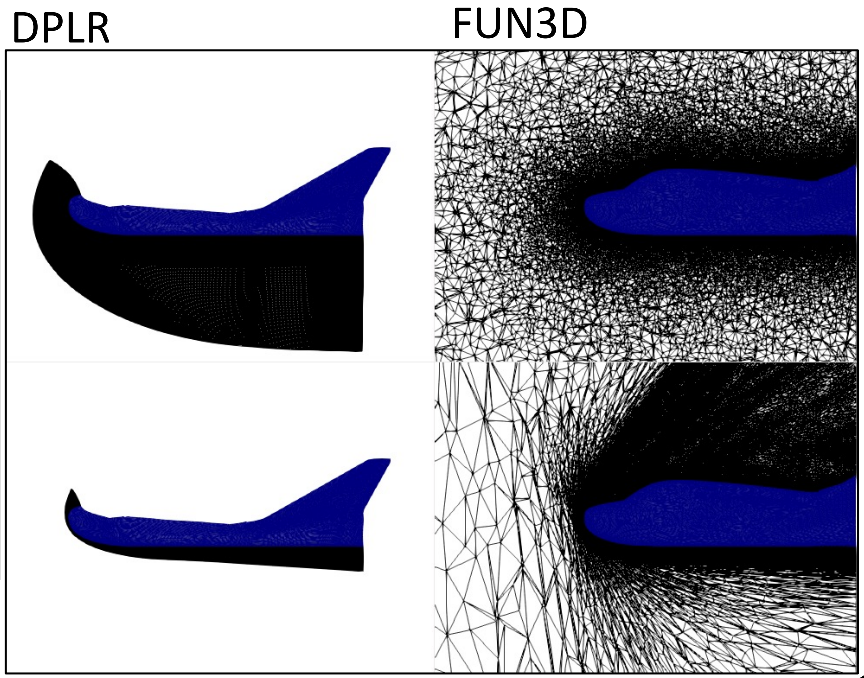

[1]Nastac, G., Walden, A., Wang, L., Nielsen, E., Liu, Y., Opgenorth, M., Orender, J. and Zubair, M., A Multi-Architecture Approach for Implicit Computational Fluid Dynamics on Unstructured Grids, AIAA 2023-1226.

Sierra Space Dream Chaser

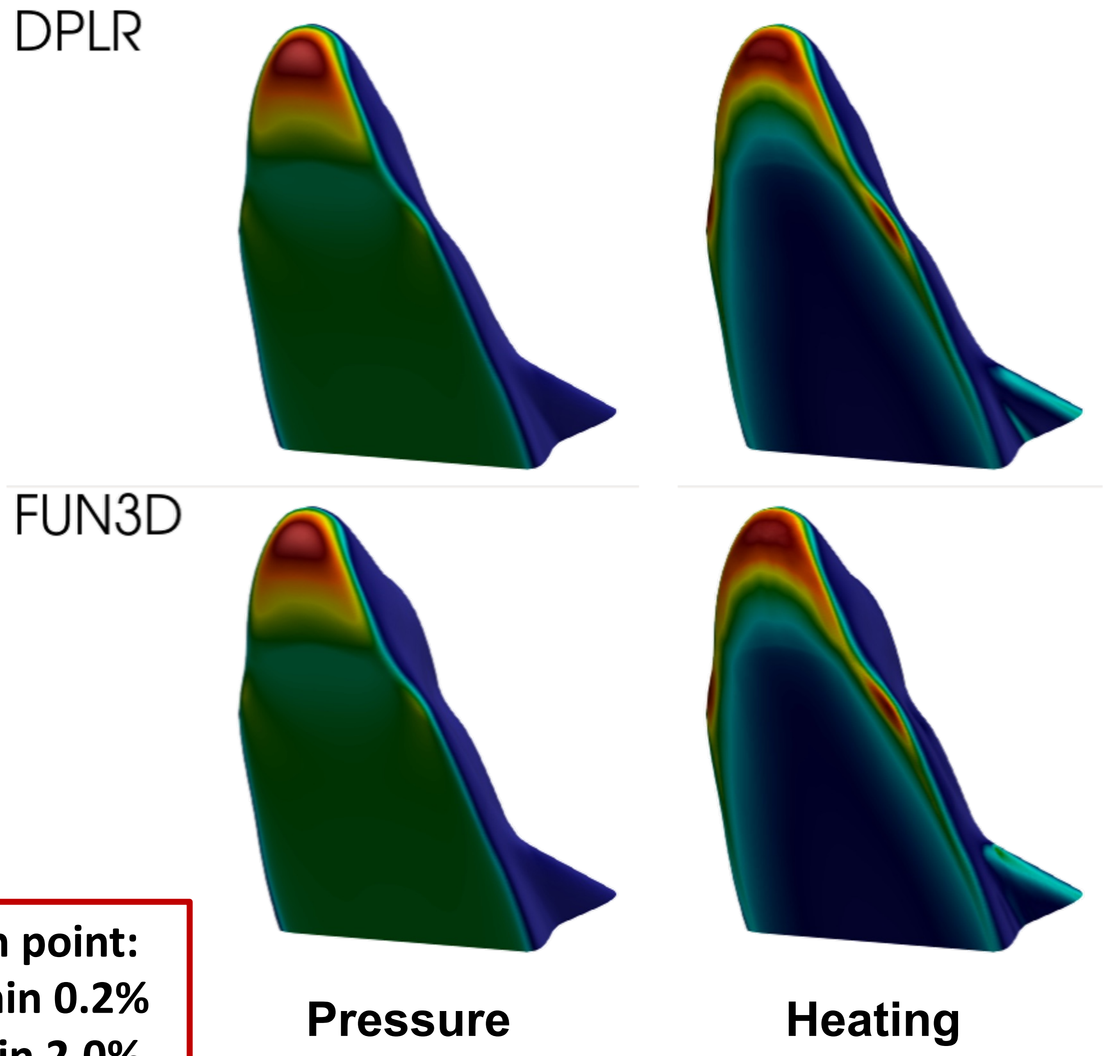

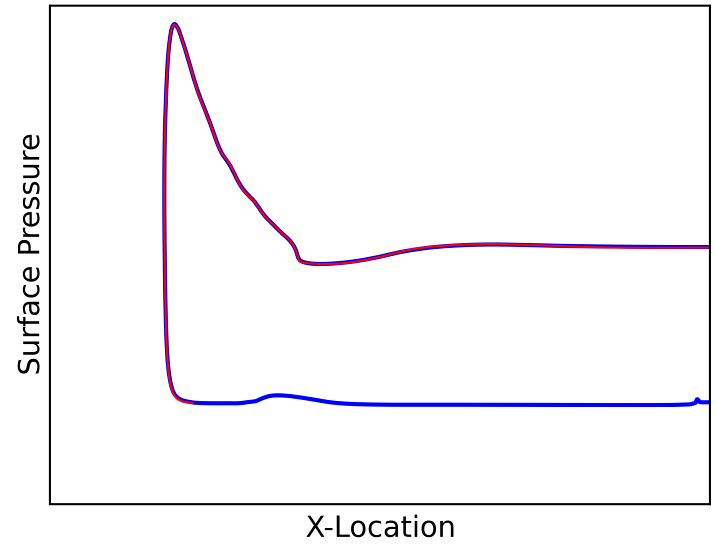

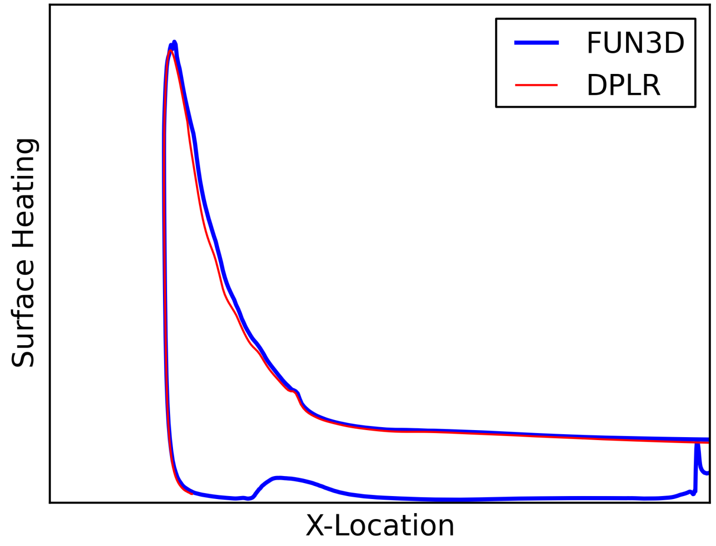

At stagnation point:
Pressure within 0.2%
Heating within 2.0%

Summary

- Enabling robust, accurate, and efficient CFD to improve next-generation aerospace vehicles
- GPUs reduce time to solution and power footprint
- FUN3D simulations demonstrated across the world's only exascale system, equivalent of O(30M) CPU cores
- Coupled FUN3D / POST2 enables CFD-based trajectory analysis
- Grid adaptation from geometry routinely used for steady-state simulations

Thanks to all of the FUN3D team, our collaborators, and partners that helped make this happen!

gabriel.c.nastac@nasa.gov

Assistance from the technical staff at Advanced Micro Devices, Inc., Intel Corporation, and NVIDIA Corporation is greatly appreciated. Support from the HPE/Cray Frontier Center of Excellence, Oak Ridge Leadership Computing Facility (OLCF) Frontier Center for Accelerated Application Readiness (CAAR) program, and the Argonne Leadership Computing Facility (ALCF) are also acknowledged. This research used resources of the OLCF at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This research also used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract No. DE-AC02-06CH11357. The authors would like to thank Christopher Stone for his contributions to AMD HIP extensions, optimizations, and the autotuning framework. The authors would like to thank Sameer Shende of the University of Oregon. The authors would also like to recognize the support of Intel Corporation through the Intel oneAPI Center of Excellence located at Old Dominion University. This research was sponsored by the NASA Langley Research Center CIF/IRAD program, and the NASA Transformational Tools and Technologies (TTT) Project of the Transformative Aeronautics Concepts Program under the Aeronautics Research Mission Directorate.